

Abstract of the Disclosure

REPLACEMENT SELECTION WITH DUPLICATE KEY HANDLING

The replacement selection algorithm for ordered merging is improved by remembering at each node of the selection tree when a losing data item is a duplicate of the prior winning data item at that node. The next time a key comparison between data items would be performed at such a node, it is not, and instead the duplicate prior loser is instantly promoted to be the next winner, avoiding the need for any further key comparisons for that winner. A data item's "duplicate" status in a node is indicated by a value of a status indicator having other values corresponding to "empty", "merging" and "done". The status indicator is an integer variable in the array data structure representing the selection tree, and has values 0 (empty), 1 (duplicate), 2 (merging), and 3 (done), so that comparisons of two integer values of status indicators can determine whether a duplicate value is ready for instant promotion to become the winning data item.